Date: 06/20/2019

EIC Detector R&D Progress Report

Project ID: eRD23
Project Name: Streaming readout for EIC detectors
Period Reported: from 01/01/2019 to 6/30/2019
Project Leaders: M. Battaglieri and J. C. Bernauer
Contact Person: M. Battaglieri and J. C. Bernauer

Project members
J. Huang, M. Purschke
Brookhaven National Laboratory, Uptown, NY

S. Ali, V. Berdnikov, T. Horn, M. Muhoza, I. Pegg, R. Trotta
Catholic University of America, Washington DC

M. Battaglieri, A. Celentano
INFN, Genova, Italy

D. K. Hasell, R. Milner
Massachusetts Institute of Technology, Cambridge, MA

J. C. Bernauer
Stony Brook University, Stony Brook, NY and Riken BNL Research Center, Uptown, NY

C. Cuevas, M. Diefenthaler, R. Ent, G. Heyes, B. Raydo, R. Yoshida
Thomas Jefferson National Accelerator Facility, Newport News, VA

Abstract

A detector for the future Electron-Ion Collider will be one of the few major collider detectors to be built from scratch in the 21st century. A truly modern EIC detector design must be complemented with an integrated, up-to-date readout scheme that supports the scientific opportunities of the machine, improves time-to-analysis, and maximizes the scientific output. A fully Streaming Read Out (SRO) design delivers on these promises, however, it can also impose limitations on the characteristics of the sensors and sub-detectors. The streaming readout consortium will research the design space by evaluating and quantifying the parameters for a variety of streaming readout implementations and their implications for sub-detectors by using on-going work on streaming readout, as well as by constructing a few targeted prototypes particularly suited for the EIC environment.

Glossary

Front-end electronics (FEE): The electronics which interfaces with the detector, typically converting the analog signal from the detector via an analog-to-digital (ADC), charge-to-digital (QDC), or time-to-digital (TDC) converter into the digital domain.

Triggered readout: A data acquisition system in which hardware produces an electrical signal according to a trigger criterion based on a subset of detector information available quickly. The signal is used to control the conversion of detector signals into the digital domain, or to trigger the read-out of a data-window from a continuously filled buffer.

Second-level / high-level trigger: In triggered systems, higher-level triggers are often used to reduce deadtime (via a fast clear) or data amount (by dropping the so-far recorded data for that event). Each level in such a system typically has different time constraints and complexity limits. For example: a certain time frame could not be forwarded to the tracker if certain conditions are not met. In certain, complex, triggered setups, the later stages can resemble a streaming system, where a stream of events flows through a network of analysis nodes, and data selection criteria either accept or drop the event. The main remaining difference for this part is then that the data is organized and tagged by an event number instead of time stamps.

Pipelined/buffered readout: A triggered readout system where event data is stored on the front ends and read out asynchronously by the backend.

Streaming readout: A data acquisition system without an element producing electrical signals to control the conversion into the digital domain or readout of a buffer. Each channel, independently, record data over a certain threshold and stream them to a CPU farm for further elaboration.

Zero suppression: Removal of data if close to the no-signal level of the detector. For example, in ADC data, removal of the pedestal.

Noise suppression: Removal of data produced by intrinsic or extrinsic detector noise, for example by correlation with neighboring channels or shape analysis.

Feature extraction: Calculation of higher-level information. E.g. calculation of hit time and energy from ADC samples, or calculation of track information from hits. Often, but not necessarily, accompanied with the removal of the underlying lower-level data.

Online Physics analysis: Analysis of the high-level information provided by the feature extraction steps to produce physics-relevant information (e.g. missing mass).

Data selection: In a SRO system, data can be algorithmically selected for further processing and long-term storage. Not selected data is dropped and not further processed. This is equivalent to the function of first and higher-level triggers in triggered systems but can make use of all detector information and results from further analysis steps including feature extraction and online physics analysis.

Streaming readout for EIC

Similar to triggered systems, which range from simple one-level synchronous designs to multi-level-trigger pipelined systems, a streaming readout solution can span a wide gamut of design complexity and trade-offs.

For example, all LHC experiments currently implement streaming readout systems on one extreme end, driven by limits on the amount of data that can be archived for long-term storage. The streaming readout systems used in these experiments contain aggressive zero suppression, noise suppression, feature extraction and physics analysis steps, dropping most if not all of the raw information and only store final physics information.

Other implementations, like the Compressed Baryonic Matter (CBM) experiment, require SRO to realize complicated, signature-based, data selection criteria only possible with tracking, but can save the underlying raw information for later reprocessing, for example with better tracking algorithms or better calibration.

On the other end are SRO system designs like the one for The Cubic Kilometre Neutrino Telescope (KM3NeT). They are not driven by data rate or trigger constraints and could be realized with a classic triggered readout system. Here, SRO is chosen mainly because it simplifies the design and reduces the risk of unrecoverable design errors.

The exact position of EIC on this scale depends on details still in flux, like detector design, maximum luminosity etc., however it is likely that the total data rate after zero-suppression and a rather limited amount of feature-extraction and noise suppression is low enough to keep all data without any selection bias (see below for a rate estimate for an sPHENIX type EIC detector). This would allow to define data selection criteria post data taking, maximizing the opportunity for data mining and searches for physics not covered by a trigger.

Further, at the rates expected at the EIC, certain detector concepts require streaming readout to be efficient. For example, the drift time in a TPC necessitates long read-out times and would incur large deadtimes in a triggered system. While a hybrid system, where the TPC electronics is streaming but fast detectors are read out in a triggered mode, are possible (see sPHENIX), they increase the overall complexity of the system with minimal or no cost reduction.

Further advantages of a SRO for an EIC include:

- Many modern FEE already are streaming capable but include a digital buffer and additional logic to mimic a traditional, triggered system. A SRO would reduce the complexity and cost of these FEE elements.
- A SRO relaxes hard timing constraints, making the overall system simpler, less error-prone and easier to debug.
- Implementation of large parts of the DAQ move from hardware into firmware and software, making the system more flexible.
- At the same time, it widens the pool of people with relevant expertise.
- A SRO removes the bottleneck of event-building from the main data flow, eliminating an error-prone component and allowing for simpler scaling of the overall system.
- SRO facilitates the convergence of offline and online analysis, allowing for better quality control online and better accelerator/experiment integration.
- Auto-calibration without trigger-introduced bias is possible.

- A SRO can continue data taking with minor degradation even if few channels are lost. In a triggered system, hung FEE boards typically stop the whole system if no additional steps are taken.
- Using CW beams or beams with extremely short bunch spacing compared to typical event durations, together with high event rates, the attribution of detector signals to specific events can be complicated. Time-tagged data allows to identify events with higher-level information, simplifying this task.
- Full validation of the complete read-out system via simulation is possible before as well as during the actual experiment by running Monte Carlo events through the same online data processing system. This allows the performance of the system, to be determined a priori, such as the efficiency, and to eventually find and fix errors and issues. In case of a traditional, triggered DAQ system, such studies are much harder, and almost always requires the implementation of a software emulation of the actual trigger behavior.
- Furthermore, a streaming readout system for the EIC can leverage the current and future advances in electronics, computation, and data storage to maximize the physics yield.

The main disadvantage of the transition to SRO is that the nuclear physics community has less experience with such systems. We therefore want to build experience through R&D on smaller test setups and experiments. A SRO also precludes the use of certain ASICs like APVs or DREAM chips which multiplex multiple analog signals into one ADC. However, typically, such chips would introduce large dead-times: at 100kHz, a 128-channel chip like the APV, read out via a 40 MHz ADC, would produce at least 30% deadtime.

Past

What was planned for this period?
What was achieved?
What was not achieved, why not, and what will be done to correct?

BNL

BNL has experience developing streaming readout (SRO) electronics. In the past decade, two major detector upgrades on RHIC detectors relied solely on streaming readout front-ends, the Forward Silicon Vertex Tracker (FVTX) at the PHENIX experiment[1] , and the electron end-cap time-of-flight (eTOF) at the STAR experiment[2]. In the next-generation RHIC detector, the sPHENIX experiment, all three tracking detectors utilize SRO capable ASICs, which continuously digitize the signals on-detector and transmit zero-suppressed data stream via optical links to “Front-End Link eXchange” (FELIX) cards, originally developed for ATLAS. They also distribute the RHIC clock to the FEE. The comprehensive system of SRO detectors in sPHENIX can serve as a foundation for streaming readout of the entire apparatus at an EIC and provide reusable R&D efforts in a realistic setting in a collider environment.
In this section, we report work that includes EIC detector rate simulation, DAQ requirements, a FELIX-based SRO DAQ concept, and prototyping progress on ASICs and DAQ. This work is funded under BNL LDRD 19-028 in synergic with sPHENIX Advanced R&D.

[1] DOI:10.1016/j.nima.2014.04.017
[2] arXiv:1609.05102

Rate study

EIC is a unique collider with diverse physics topics, which imposes unique requirements on the DAQ design. Comparing to hadron colliders, the EIC has much higher crossing frequency (100-500 MHz) and higher luminosity (10^{34} cm^{-2} s^{-1}). However, the EIC cross section (~40 μb) is much smaller than the proton-proton cross section (40-80 mb from RHIC to LHC), making the EIC collision rate (~500 kHz) and particle production rate (4M charged particle per second) a factor 10-1000 lower than that of RHIC and the LHC. Nonetheless, the EIC has a diverse event topology and is sensitive to the machine and detector background. All these effects pose challenges to EIC DAQ design, and its requirements should be based on realistic detector simulations.

In this subsection, we study the EIC detector data rate using full detector Geant4 simulations (see Figure 1), based on the sPHENIX-based EIC detector concept[3]. Such estimates would also apply to many other EIC detector concepts. Detailed rate simulations are described in the 2018 EIC detector design study report for the sPHENIX-based EIC detector[4]. A follow-up study was carried out to estimate beam-gas interactions and its induced background rate.

The data rate from tracking and calorimetric detectors are summarized in Figure 2. The overall rate with PID detector and moderate detector noise would be around 100 Gbps for all EIC collisions. Put into context, the total signal data rate fits well within the designed DAQ disk-writing bandwidth for the sPHENIX DAQ, which is at least 200 Gbps [sPHENIX TDR].

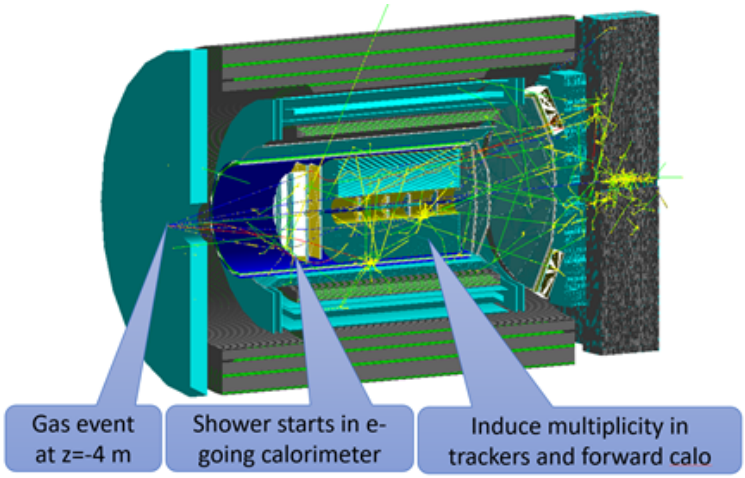

Figure 1: Geant4 simulation of a beam gas interaction background in the sPHENIX-EIC detector concept, which is the basis for the overall detector rate estimates and defines the SRO DAQ throughput requirements.

[3] arXiv:1402.1209

[4] sPH-cQCD-2018-001, publicly available at https://indico.bnl.gov/event/5283/

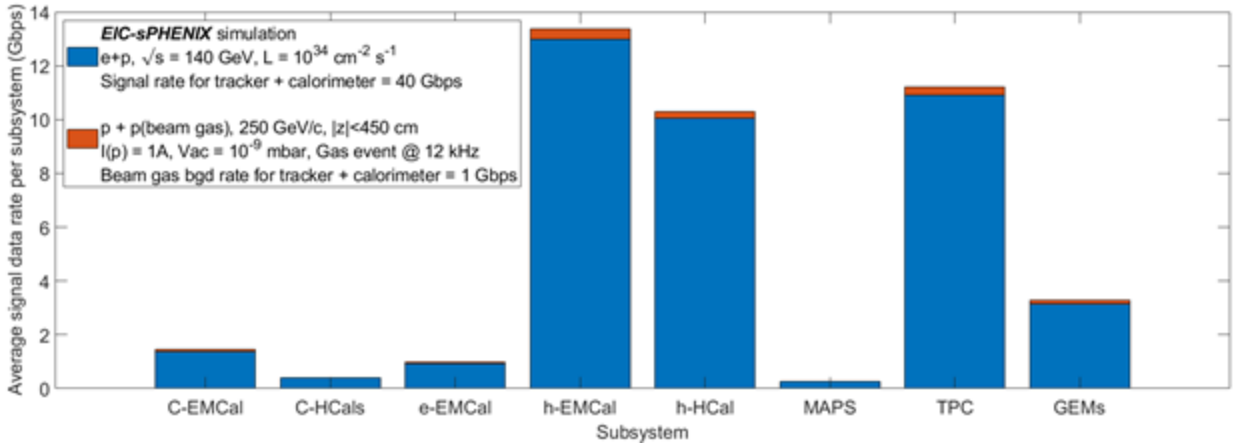

Figure 2: Signal data rates from tracking and calorimetric detectors from EIC collisions (blue) and beam gas interaction (red) via full detector Geant4 simulation of the sPHENIX-EIC detector concept [sPH-cQCD-2018-001]. Zero-suppression and realistic data format based on sPHENIX prototyping are assumed in this estimation. The overall tracker data rate is 40 Gbps. The estimated rate with PID detector and moderate detector noise would reach 100 Gbps for full experiment.

Rate & Physics Driven SRO DAQ Requirements

The estimated data rate allows for long term storage of almost all data with minimal reduction, and a SRO solution would secure all the advantages listed above. The following list summarizes the requirements for an EIC DAQ

- Continuous digitization of the detector signals by all front end electronics
- Reliably synchronization of all front-ends nominally to the of the beam crossing clock, and identification of rare faults
- Recording of the collision data stream (100 Gbps in raw data)
- If needed, filtering out background/noise with minimal signal loss
- Reliable data flow and control systematics

FELIX-based SRO DAQ

In response to the EIC DAQ challenges, we have developed a SRO DAQ concept for EIC based on the BNL-designed FELIX card (Figure 3), which is used to bridge the custom front-end to the hosting server (EBDC). The FELIX card carries a large FPGA (Xilinx Kintex UltraScale KU115) and supports 48 bi-directional, 10-Gbps capable optical links to the FEE and a 100-Gbps PCI-express Gen3 link with the hosting server's CPU. A mezzanine card is used to receive sPHENIX timing and trigger information. The next generation will support 25Gbps optical transceivers and PCIe Gen4 interface. This DAQ interface was initially developed for the ATLAS Phase-1 upgrade and beyond. Similar architectures are also used in many DAQ systems in the 2020s, such as the upgrades of LHCb, ALICE, the sPHENIX and CBM experiments.

In this DAQ concept (Figure 4), experiment timing and slow control are directly distributed to FEEs via FELIX optical links to FEEs. All FEEs are synchronized with the beam crossing clock and checked via a synchronization signal. The data stream is sent from the FEE to the FELIX with zero-suppressions applied either on FEE or on FELIX, which supports 500Gbps data input

bandwidth for each FELIX card. Further noise filtering could be performed on the FELIX FPGA. Then raw data is buffered on the commodity host server prior to sending to network storage servers. The FELIX-server buffering can also bridge the custom detector FEE with μs buffering with the commodity network with seconds of data buffer depth. The event reconstruction from the data stream is performed offline, allowing time for final calibrations and opportunities of data mining with specialized reconstructions.

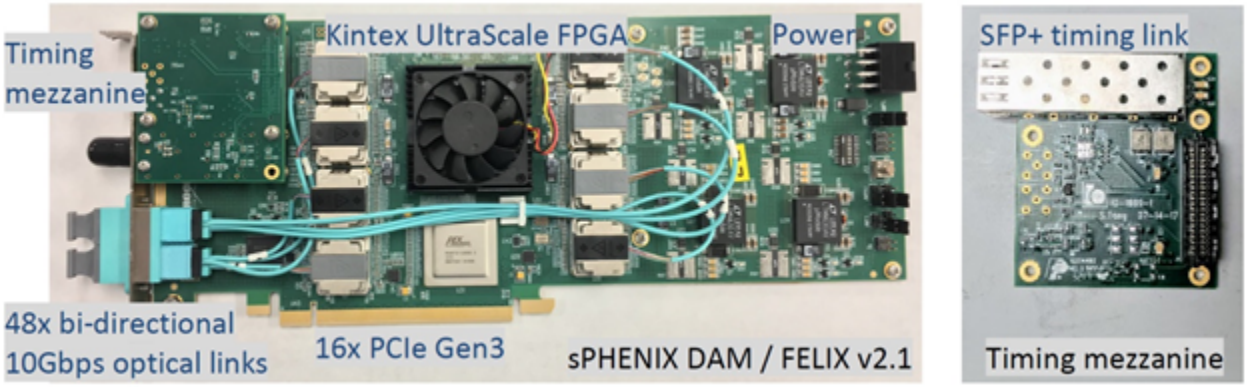

Figure 3: FELIX v2.1 DAQ interface card, which is a PCIe card providing 500 Gbps bi-directional optical links to FEEs and a 100-Gbps PCI-express Gen3 16-lan link with the hosting server. This DAQ interface is developed at BNL, initially for the ATLAS experiment, and now also adopted in the Proto-DUNE, sPHENIX and CBM experiments.

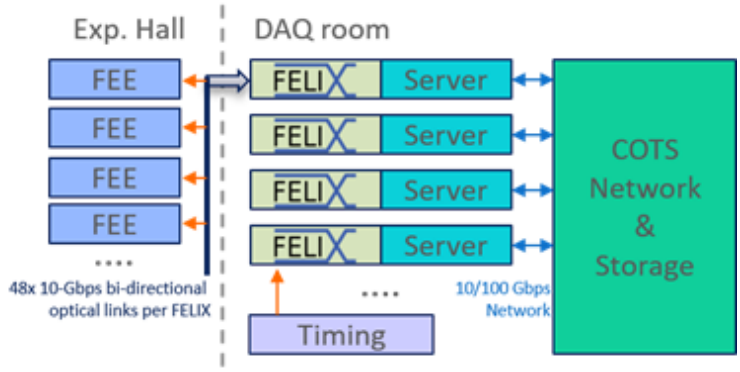

Figure 4: FELIX-based EIC DAQ concept. All FEEs are synchronized to the beam interaction clock via the FELIX interface. The digitized data are streamed and buffered on the FELIX FPGA and server prior send to network storage servers in counting house. This is very similar to the DAQ setup for the sPHENIX tracking detector systems.

Prototyping of streaming ASIC and DAQ

Significant progress has been made on prototyping SRO ASICs in conjunction with sPHENIX advanced R&D. Two ASICs are highlighted in this report.

- The SAMPA ASIC was initially developed for the ALICE upgrade [ALICE TDR], integrating 32-channel shaping, 10bit SAR ADC, a DSP function supporting both triggered and SRO modes, as well as on-chip zero-suppression. It can be applied to many tracking detector applications in the EIC era, including TPC and GEM tracking detectors. A current test-stand is shown in Figure 5, with an 8-SAMPA FEE board used to read out eRD6 HBD-

TPC prototype via a FELIX card. A beam test of this readout system is ongoing at Fermilab Test Beam Facility in June 2019. The next version of the SAMPA chip is also under development at the University of São Paulo under sPHENIX R&D optimized for 80ns shaping time and 20MHz waveform digitization.

- The FPHX ASIC was initially developed for the PHENIX FVTX tracker[5], integrating 128 channel shaping, 3-bit Flash ADC, and digital buffers. It operates solely in an SRO mode. New silicon ladders read out with this ASIC (see Figure 6) are being tested in June 2019, to demonstrate the operation via 1.3-m bus extension. Such a long extension will be critical for EIC application of silicon tracker by facilitating the integration of silicon tracks in the tight but material sensitive space in the vicinity of the interaction point.

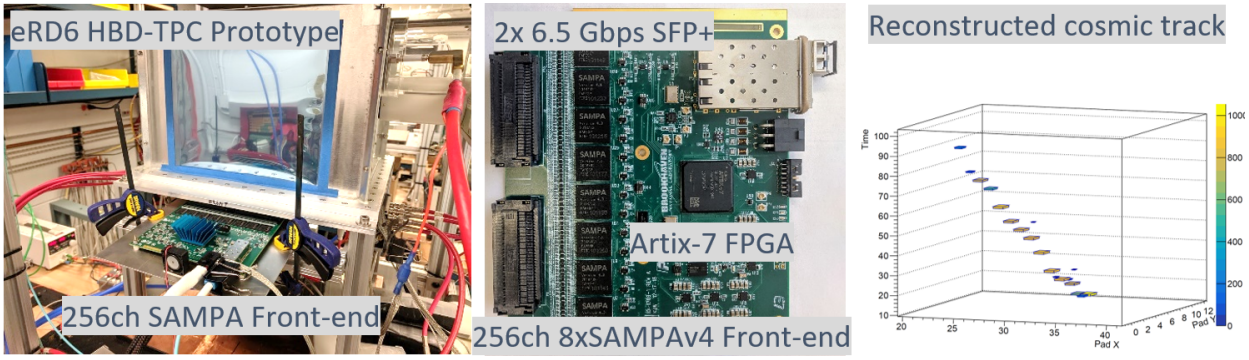

Figure 5: 8-SAMPA FEE board (middle) used to readout eRD6 HBD-TPC prototype (left) via FELIX v2.1 at BNL. The data was reconstructed offline showing a cosmic track as shown on the right panel.

Figure 6: Left: silicon strip tracker with FPHX wire bonded to the sensor, developed by RIKEN and BNL teams for sPHENIX advanced R&D. Four layers ladders containing 3k channel each are being tested at Fermilab Test Beam Facility in June 2019. Right: PHENIX/sPHENIX readout board bridging LVDS signal from FPHX to optical links. The recent test demonstrated its connection with FELIX interface.

[5] DOI:10.1016/j.nima.2014.04.017

INFN-GE/CUA

During the period August 2018-July 2019 we setup a fully working triggerless DAQ system, based on the Waveboard DAQ board for SiPMs developed at INFN and on the TriDAS readout system from the KM3NeT experiment. The system has been used to readout the Beam Dump eXperiment "BDX-mini" detector currently installed at JLab, comparing the performances with those obtained from a traditional triggered readout.
The Waveboard hosts 12 input channels for SiPM, each providing analogue amplification, sensor bias voltage, and digitization. The board operates in streaming-mode: whenever the input signal crosses a programmable threshold, digitized samples within a programmable window are streamed through the network via TCP to an online CPU processing farm. Each waveboard can be synchronized to a GPS system, in order to label each sample with an absolute timestamp that allows for channel-to-channel matching in the online system. The absolute accuracy of the timestamp is better than 1 ns. The collection of digitized samples, plus with the associated time-stamp, defines a "hit".
The online reconstruction system TriDAS receives a continuous stream of hits from individual detector channels. These are ordered in time, and then divided into independent time-windows, here referred to as "time-slices". Each "time-slice" is then analyzed by a CPU farm node, to identify and store events there present. The system can be easily scaled in case in case more computing power is needed to analyze a time-slice, by simply adding further nodes in the online farm.
The analysis of each time-slice proceeds as a two-step process:

- In the first step, a "L1 algorithm" scans the full time-slice and identifies "events". An event is defined as a collection of hits within a programmable time-window (of about few us) centered on a "L1 seed hit". A "L1 seed hit" is here defined as a detector hit - from any channel - with an integrated waveform area exceeding a programmable threshold.
- In the second step, one or more "L2 algorithms" are run on each event within the waveform. Each L2 defines an event selection algorithm, tailored to a specific physics or calibration channel. For each L2 algorithm, each event is marked with a Boolean flag reporting whenever that event satisfied the corresponding criterium or not.

The TDAQ system was installed in March 2019 at Jefferson Lab and used for BDX-mini detector readout. BDX-mini is an array of two $PbWO_4$ crystal matrixes, each made by 22 samples, surrounded by two layers of plastic-scintillator counters. Both crystals and plastic scintillators are read by SiPM. The detector is currently deployed at Jefferson Lab, in a vertical pipe installed behind the Hall-A beam dump, placed at the nominal beam height.
The detector is currently used in an experimental measurement campaign conducted by the BDX collaboration to explore light dark matter (LDM) produced in the interaction of the CEBAF beam with the Hall-A beam dump. After production, LDM would travel through the Beam-dump shielding and following dirt, finally reaching the detector and possibly interacting with the $PbWO_4$ crystals, with a visible energy deposition. Due to the current low-energy (2.2 GeV) CEBAF operations, no beam-related backgrounds are foreseen, and the only expected signals are those from cosmogenic backgrounds.
The results obtained running the BDX-mini in triggerless mode have been compared with those obtained with a traditional triggered system, finding excellent agreement and thus **validating this technology for the small-scale, low-rate system that we considered.** In particular, the following tests have been performed:

- The differential energy spectrum for individual crystals and the total energy spectrum - summed over different channels - have been compared, finding an excellent agreement.

- The effect of a L2 “clustering” algorithm has been checked, finding an efficiency of approximately 100% for the events above the programmable cluster threshold.

JLAB: SAMPA / TPC streaming

Streaming Readout for a two plane Gas Electron Multiplier[GEM] detector was planned for JLAB and hardware was acquired to instrument two GEM layers that have 384 channels each. The Front End Card [FEC] is from ALICE, and has modifications made at JLAB. Each of the FEC cards has 5 SAMPA ASIC, so each card manages 160 GEM channels. There are five FEC cards, so there are plenty of channels to fully map the two GEM layers.
Each of the FEC cards interfaces to a ReadOut-Receiver Card that was designed for the ALICE/ATLAS experiment. These RORC cards interface with the Front-End Cards by a use of the CERN GigaBit-Transceiver serialization protocol. The fiber optic link connects each of the five FEC with the RORC card. The RORC is installed in a PC chassis that uses PCIe to transfer data to the CPU.
Data is continuously streaming from the FECs to the RORC. Presently, a simple trigger that consists of two scintillators that are placed above and below the two GEM layers, create the signal that initiates data transmission of a programmable window of streamed data to be captured by the RORC. This data is transmitted to the PC memory and written to disk.
Most of the 9 Gb/s data sent from each FEC to the T-RORC (45 Gb/s total) consists of sync packets that serve to keep the serial links from the SAMPAs active when there is no hit data to send. With the current firmware the T-RORC would try to transmit all of this data directly to PC memory. We use the triggering scheme to keep data volume to memory and to disk at a manageable level.
To truly acquire data in a continuous fashion, there will need to be modifications of the RORC firmware to suppress the transmission of unnecessary sync packet data to memory and disk. This development is presently ongoing and will be completed this summer.

MIT

In principle MIT intended to study the front-end electronics in collaboration with the ASIC chip design and manufacturing firm Alphacore. The plan was to match various detector types (GEM, PMT, SiPM, etc.) with ASIC chips with various sampling frequencies, number of bits, input capacitances, etc. to develop a catalogue of front-end electronic evaluation boards that could be paired to FPGA boards for testing. Unfortunately, we are still waiting on budgetary approval for 2018-2019 and thus have not had the funds to order ASICs or to build simple prototype detectors.

However, with the help of Tanja Horn at Catholic University of America we are borrowing a small number of lead tungstate crystals that will be assembled into a calorimeter for testing in September at the DESY test beam facility. We hope to read this out using both a traditional readout scheme and a streaming readout scheme. The results of these tests will study not only a streaming readout scheme but also evaluate the feasibility of using lead tungstate calorimeters at future EIC experiments and a possible two-photon exchange experiment.

SBU/RBRC

Together with M. Diefenthaler and D. Romanov, a requirements catalog and first concept for a detector agnostic, universal on-wire data format has been designed. Prototype implementation and further discussion at the Streaming Readout Workshop IV has led to further improvements. The requirements for the protocol can be summarized as follows:

- Universality: The format must allow encapsulation of arbitrary detector data. The format of that detector data cannot be foreseen and should not be fixed by the standard.
- Scaling: The format should scale from a single-channel table-top experiment to millions of channels. This secures that personnel can be trained at small installations at their home institute as well as that the format can be used pre-EIC for detector tests etc. to gain experience.
- Error resilience: It should be possible to detect and recover from partial corruption of the data.
- Bandwidth efficiency: To save bandwidth in the read-out network and on disk, the format should incur a minimal overhead.
- Efficiency for merge/drop/seek: The format should make it possible to merge different streams sorted by a time tag without having to parse the full information, ideally allowing zero-copy optimizations for merge-type operations and minimal overhead if data is deselected by filter operations. Further, it should be possible to seek forward in the stream without having to read all data.
- Ease of implementation in software: The protocol should be easy to implement on CPUs and language agnostic so that bindings to many software languages are possible. Not selecting a canonical language will future-proof the standard. For example, there is currently a strong movement away from classical languages like C and C++ to more high-level languages like Python or languages making it easier to write safe and parallel code like Rust and Go.
- Ease of implementation on FPGAs: At the same time, the protocol must be straight-forward to implement on FPGAs with good performance and minimal resource usage.

Established on-disk formats like ROOT files or HDF5 are optimized for disk storage and CPUs. They are therefore not well suited for this purpose.

The current protocol draft assumes that the data is packetized by time-frames. Each packet consists of these data fields:

1. Constant start byte (0xAA) to resynchronize reading the data stream in case of corruption.
2. 24-bit sequence number to reorder packets. Since packets can be processed by different nodes in a parallel processing model, and processing time is variable, packets can arrive out-of-order at later stages. A sequence number allows to efficiently reorder the packets.
3. 32-bit channel id. Every unique data stream source, i.e. FEE modules and processing nodes, mark the produces channel with a unique id.
4. 32-bit block size, as the remaining 32-bit words of this packet. This allows to easily seek to the next packet.
5. 16-bit flags. So far, none of the flag bits are defined.
6. 48-bit coarse frame time counter. Together with a 32-bit fine time at a 1 ps resolution, the system would have 4ms time-frames and can encode a time span of 38000 years. Ideally, the times are absolute, that is, in reference to a certain fixed date, instead of relative to run start or similar.
7. 32-bit aligned semi-opaque payload data. The payload data can be up to block-size words long. If more space for the data of a time-frame is required, multiple packets with

the same coarse time need to be used. The detector data is undefined; however, the payload data is required to be an array with this structure:
    a. 32-bit fine time
    b. 32-bit aligned detector “hit” data

    With this scheme, a common middleware can be written that can present data on a time-frame or individual hit level to the processing algorithms, without prior knowledge of the detector information data structure.
8. 32-bit CRC checksum to detect on-wire and on-disk corruption. While on-wire corruption should be detected by the underlying transport, on-disk corruption is often silent and needs to be detected at the protocol level. We do not plan to use error correction codes (ECC): the loss of single time frames is acceptable, while ECC had complexity and run-time costs.

The whole data structure is 32bit aligned with little-endian byte ordering, minimizing data manipulation on the prevalent CPU architectures.
Information like channel names and their mapping to IDs, the size of the detector hit-data, etc. is not saved in the data structure. Further infrastructure must hold a database of these values. We currently research if established internet standard protocols could be used, for example, channel names to IDs as well as further channel information could be stored via DNS.

We started to discuss overlap with the software consortium. A standard protocol could also be used for the data exchange in MC and other physics studies. Another direct point of overlap are MC implementations with a streaming, that is, with a timestamp-oriented instead of event-oriented data organization design, which would easily allow realistic pile-up studies, and much of the work on these studies already goes in that direction.

sPHENIX will be the first large scale nuclear physics experiment to use streaming readout for a main part of the detector. The group is involved in the DAQ software realization.

Further, the SBU/RBRC group is involved with the DAQ system for the calorimeter test beam at DESY described above.

Future

What is planned for the next funding cycle and beyond? How, if at all, is this planning different from the original plan?
What are critical issues?

BNL: BNL plans to produce multiple FELIX v2.1 for use in support of the EIC SRO DAQ prototyping. ASIC development will proceed in the next generation of SAPMA, FPHX and ALPIDE chip and their integration in a SRO tracking system in synergic to the sPHENIX advanced R&D. We will also work towards a full system testing integrating EIC timing distribution, SRO trackers, and calorimeters.

INFN-GE/CUA: In the forthcoming funding cycle we plan to perform a test of our triggerless TDAQ system with a calorimeter exposed to a test-beam. A discussion is currently ongoing between INFN-GE/CUA and MIT to identify a possible detector and test location. The goal of the test is to validate the streaming readout technology within a high-rate application, where the requirements to the acquisition boards and to the online CPU farm are more demanding.

SBU/RBRC: We plan to write an RFC-style documentation of the protocol. The standard will then be implemented independently at JLAB and SBU, followed by an interoperability test. The protocol will be used in the upcoming calorimeter test beamtime at DESY, in collaboration with the MIT group. Further, we will investigate solutions for other parts of the required software framework, including network protocols, device bring-up, channel-mapping and node orchestration.

JLAB: The development of a streaming TDC that will be prototyped on a standalone format is planned as an extension of the ongoing work that has been completed for studying the clock and synchronization distribution for a large streaming readout system. The TDC prototype will have 64 channels, with 20ps resolution and include standard PCIe_Gen3 and 10Gb Ethernet interfaces. The 10Gbe connections will be tested with our Xilinx KCU1500 Accelerator boards in conjunction with the new Arista 7170-32 network switch to demonstrate data rates and provide a platform for software development that will ultimately handle the streaming data.

The firmware development for streaming digitized data from the JLAB FADC250 through the VXS Trigger Processor[VTP] is nearly complete and ready for testing with the ZeroMQ messaging system. The performance of the TCP/IP firmware has been simulated at an Ethernet rate of 9.4Gbps so it fits within the 10GbE bandwidth. This new firmware combined with the ZeroMQ data messaging wrapper will be tested before the end of FY19.

MIT: As previously stated our effort in the streaming readout consortium will be to work with ASIC chip design and manufacturers to develop an integrated catalogue of evaluation boards that can be easily implemented as the front-end electronics for a variety of standard nuclear and particle detectors. The catalogue would include a range of options for sampling rate, number of bits, input capacitance, etc. so that a future user could simply choose an evaluation board based on their experimental needs. The evaluation boards would also be paired to standard FPGA boards that could be programmed to perform zero suppression, feature extraction, timing, and then streaming the results to a CPU cluster as needed for further analysis or storage.

In addition to this work we will continue our efforts to study lead tungstate calorimeters at the DESY test beam. These studies will not only evaluate the possible future use of such crystals in an EIC detector but will also be used as a test bed for comparing traditional readout schemes with a streaming readout approach.

Manpower

Include a list of the existing manpower and what approximate fraction each has spent on the project. If students and/or postdocs were funded through the R&D, please state where they were located, what fraction of their time they spend on EIC R&D, and who supervised their work.

All personnel are currently funded by external sources. We report here time spend on SRO related activities, whether directly EIC-related or not.

INFN-GE/CUA: the personnel involved in the aforementioned activities at INFN-Genova are: Marco Battaglieri (senior staff scientist), Andrea Celentano (staff scientist), Luca Marsicano (PhD student). Each of us has spent approximately 20-30% of the time on this activity, partially shared on synergistic activities, in particular the BDX experiment at Jefferson Laboratory.

SBU: The work on protocol and architecture is carried out by J. C. Bernauer (30%), with frequent input from M. Diefenthaler, C. Cuevas and M. Purschke.

BNL: J. Huang, M. Purschke will commit 10-20% time developing the SRO system for EIC, which will be supported under BNL LDRD 19-028 and in synergy with on-going work on sPHENIX SRO tracking system. This work will be supported by an experienced engineering team at BNL including J. Kuczewski, J. Mead, and A. Dellapenna and in collaboration with the ATLAS DAQ team at BNL.

JLAB: Streaming readout work is performed by staff scientists and engineers from the JLAB DAq and Electronics groups. G. Hayes and C. Cuevas are the respective group leaders with B. Raydo, E. Jastrzembski and J. Gu working 10% of their time on various streaming readout activities. The streaming readout development activities are an ongoing extension of the hardware developed and implemented for the 12GeV experimental programs at JLAB. These activities are collaborative with the SRO consortium groups and supported by JLAB LDRD-1910.

MIT: Main effort will be through D. Hasell, R. Milner, I. Friscic, S. Lee, P. Moran, and B. Johnston working on the DESY test beam. I. Friscic will assume a shared role at JLab and MIT working on streaming readout in the fall.

External Funding

Describe what external funding was obtained, if any. The report must clarify what has been accomplished with the EIC R&D funds and what came as a contribution from potential collaborators.

All activities described above were funded by external sources. Many of the results have been presented at the Streaming Readout Workshop IV, held in Camogli from May 22th to May 24th, which was supported by the eRD23 funds.

Publications

Please provide a list of publications coming out of the R&D effort.

Funding request

A critical issue is to identify possible streaming readout hardware solutions for each detector component, as well as to study and solve potential problems related to the intersection between SRO and detector. Therefore, we plan to organize meetings and mini-workshop with colleagues from the detector research initiatives, especially those not already represented at our regular meetings. Cognizant of the limited funds available, we only request travel funds to support these and our regular meetings as well as outreach to the EIC UG, for a total of $20,000.